

# Self-Avoiding Walks over Adaptive Triangular Grids\*

Gerd Heber<sup>†</sup>

Rupak Biswas<sup>‡</sup>

Guang R. Gao<sup>§</sup>

## Abstract

Space-filling curves is a popular approach based on a geometric embedding for linearizing computational meshes. We present a new  $O(n \log n)$  combinatorial algorithm for constructing a self-avoiding walk through a two-dimensional mesh containing  $n$  triangles. We show that for hierarchical adaptive meshes, the algorithm can be locally adapted and easily parallelized by taking advantage of the regularity of the refinement rules. The proposed approach should be very useful in the runtime partitioning and load balancing of adaptive unstructured grids.

## 1 Introduction

Advances in adaptive software and methodology notwithstanding, parallel computational strategies will be an essential ingredient in solving complex, real-life problems. However, parallel computers are easily programmed with regular data structures; so the development of efficient parallel adaptive algorithms for unstructured grids poses a serious challenge. An efficient parallelization of these unstructured adaptive methods is rather difficult, primarily due to the load imbalance created by the dynamically-changing nonuniform grid. Nonetheless, it is generally believed that unstructured adaptive-grid techniques will constitute a significant fraction of future high-performance supercomputing. Mesh adaptation and dynamic load balancing must therefore be accomplished rapidly and efficiently, so as not to cause a significant overhead to the numerical simulation.

Serialization techniques play an important role when using a finite element method (FEM) over adaptive unstructured grids. A numbering of the unknowns allows for a matrix-vector notation of the underlying algebraic equations. Special numbering techniques (Cuthill-McKee, frontal methods) have been developed to optimize memory usage and locality of the algorithms. In addition, the runtime support for decomposing dynamic adaptive grids is often based on a linear representation of the grid hierarchy in the form of a *space-filling curve*. Several researchers have demonstrated the successful application of techniques based on space-filling curves to N-body simulations, graph partitioning, and other graph-related problems.

The general idea of a space-filling curve is a serialization (or linearization) of visiting points in a higher-dimensional space. A standard method for the construction is to

---

\*The work of the first and third authors was partially supported by NSF under Grant Numbers NSF-CISE-9726388 and NSF-MIP-9707125. The work of the second author was supported by NASA under Contract Number NAS 2-14303 with MRJ Technology Solutions.

<sup>†</sup>Visiting Scientist, Computer Architecture and Parallel System Laboratory, University of Delaware, Newark, DE.

<sup>‡</sup>Senior Research Scientist, MRJ Technology Solutions, NASA Ames Research Center, Moffett Field, CA.

<sup>§</sup>Associate Professor, Department of Electrical and Computer Engineering, University of Delaware, Newark, DE.

embed the object under study into a regular environment (where the standard space-filling curves live). However, this approach introduces an artificial structure in that the entire construction depends on the embedding. Furthermore, this type of a linear representation ignores the combinatorial structure of the mesh which drives the formulation of operators between the finite element spaces. The question arises whether one could reduce the different requirements for a serialization in an adaptive FEM over unstructured grids to a common denominator.

In this paper, we present a new approach to constructing a *self-avoiding walk* through a triangular mesh. Unlike space-filling curves which are based on a geometric embedding, our algorithm is combinatorial since it uses the mesh connectivity only. We also show how the concept can be easily modified for adaptive grids (that are generated in a hierarchical manner based on a set of simple rules) and made amenable for efficient parallelization (using a novel indexing scheme).

## 2 Self-Avoiding Walks

A self-avoiding walk (SAW) over an arbitrary unstructured two-dimensional mesh  $\mathcal{M}$  is an enumeration of all the triangles of  $\mathcal{M}$  such that two successive triangles share an edge or a vertex. Note that a SAW visits each triangle exactly once, entering it over an edge or a vertex, and exiting over another edge or vertex. In the cases where the SAW jumps over vertices, we imply that the triangles following one another in the enumeration do not share an edge. As simple examples show, the requirement that triangles following one another in a SAW must share an edge is too strong. In general, SAWs do not exist under such conditions.

SAWs can be used to improve the parallel efficiency of several irregular algorithms, in particular issues related to locality and load balancing. However, we need special classes of SAWs that will facilitate dynamic load balancing with good locality in terms of runtime memory access and interprocessor communication. We thus consider a special class of SAWs called *proper self-avoiding walks* (PSAW) where jumping twice over the same vertex is forbidden.

We can prove that a PSAW exists for any arbitrary triangular mesh  $\mathcal{M}$ . The proof is inductive over the number of triangles in the mesh and extends an existing PSAW over to a larger mesh. The proof actually provides a set of elementary rules which can be used to construct a PSAW: starting from any random triangle of  $\mathcal{M}$ , choose new triangles sharing an edge with the current submesh and extend the existing PSAW over the new triangle. Details of the proof and the algorithm are given in [3]. In the process of deriving the existence of PSAWs, a natural mathematical framework emerges which can be easily adapted for other special classes of SAWs.

The complexity of this algorithm is  $O(n \log n)$ , where  $n$  is the number of triangles in the mesh. This complexity is obtained as follows. Since the triangles are organized in a red-black tree, a search requires  $O(\log n)$  time (the edge-triangle incidence relation can also be organized in a red-black tree). Since a search has to be performed for each triangle in the mesh, the overall complexity of our algorithm is  $O(n \log n)$ .

We have implemented a sequential algorithm for the generation of a PSAW using the `set` and `map` containers from the C++ Standard Template Library. Both containers are implemented with red-black trees; however, one may also think about an implementation using hash tables. Results indicate the dependence of the runtime on the number of triangles (or edges). According to our description of the algorithm, one would suspect

that the runtime may also depend on the triangle/edge ratio. An investigation of this possibility clearly shows that the dependence of the execution time on the triangle/edge ratio is negligible and confirms our complexity estimate for the algorithm.

A standard parallelization of the deterministic construction algorithm would be based on the observation that if the extensions of an existing PSAW did not interfere with one another, then they could be done in parallel. Such an approach is feasible in dealing with single meshes when no additional structural information is available, or with families of meshes without any hierarchical structure, although scalability would be a major concern.

### 3 Adaptive Triangular Grids

As is evident from the discussion in the previous section, a parallel algorithm for constructing a PSAW is non-trivial, and the walk has to be completely rebuilt after each mesh adaptation. However, PSAWs provide enough flexibility to design a *local* algorithm for *hierarchical* adaptive changes in an unstructured mesh. In contrast to a general adaptation scheme that produces just another unstructured mesh, a hierarchical strategy pays particular attention to the initial triangulation and uses a set of simple coarsening and refinement rules:

- When a triangle is refined, it is either subdivided into two (green) or four (red) smaller triangles (Fig. 1).
- All triangles with exactly two bisected edges have their third edge also bisected; thus, such triangles are red.
- A green triangle cannot be further subdivided; instead, the previous subdivision is discarded and red subdivision is first applied to the ancestor triangle (Fig. 2).

Such a policy also assures that, with repeated refinement, the quality of the adapted mesh will not deteriorate drastically from that of the initial mesh. This strategy is similar to that described in [1] for tetrahedral meshes.




FIG. 1. Green (left) and red (right) refinement of a triangle.




FIG. 2. Special policy for refining a green triangle.

Such a scheme exploits the structure of the adaptation hierarchy to simplify the mesh labeling process. Our goal is to modify (or adapt) the labeling only in regions where the mesh has been altered but not to rebuild the entire indexing. This indexing idea has been extensively developed in [2]. It is a combination of coarse and local schemes. The coarse

scheme labels the objects of the initial mesh so that the incidence relations can be easily derived from the labels. The local scheme exploits the regularity and the finiteness of the refinement rules to produce names for the mesh objects at subsequent refinement levels. The coarse and local schemes are combined by taking the union of the Cartesian products of the coarse mesh objects with their corresponding local schemes. Ambiguities are resolved by using a normal form of the index.

Our basic strategy for such hierarchical adaptive meshes is to apply the general algorithm only once to the initial mesh to construct a coarse PSAW, and then to reuse the walk after each adaptation. This requirement leads us to the concept of *constrained proper self-avoiding walks* (CPSAW) that exploits the regularity of the refinement rules. In a way, a PSAW is global since it is related to the initial coarse mesh, while a CPSAW is local since it is restricted to a triangle of the coarse mesh. A PSAW leaves a trace or footprint on each coarse triangle which is then translated into a boundary value problem for extending the walk over the adapted mesh. In addition, the extensions for different coarse triangles can be decoupled; thus, the construction process of a CPSAW is inherently parallel.

Figure 3 depicts an example of how a CPSAW is generated from an existing PSAW when the underlying mesh is adapted. In the left portion of Fig. 3, we show a section of a PSAW over four triangles of an initial coarse mesh. The entry and exit points are respectively marked as  $\alpha$  and  $\beta$ . These four triangles are then refined as depicted in the right portion of Fig. 3, that also shows the extension of the original PSAW to the higher level of refinement.




FIG. 3. Extending a PSAW (left) to a CPSAW over an adapted mesh (right).

The regularity of the local picture allows us to prove the existence of CPSAWs; full details are given in [3]. Moreover, CPSAWs are also appropriate for multigrid methods since they “know” the refinement level of the adapted mesh. The process of generating a CPSAW from an initial PSAW after mesh adaptation is embarrassingly parallel, since each processor is responsible for only its own chunk of the walk. However, load balancing is still an important issue that must be addressed when the coarse mesh is locally adapted.

## 4 Conclusions

In this paper, we have developed a theoretical basis for the study of self-avoiding walks over two-dimensional adaptive unstructured grids. We described an  $O(n \log n)$  algorithm for the construction of proper self-avoiding walks over arbitrary triangular meshes and reported results of a sequential implementation. We discussed parallelization issues and suggested a significant improvement for hierarchical adaptive unstructured grids. For this situation, we have also proved the existence of constrained proper self-avoiding walks.

The algorithms presented allow a straightforward generalization to tetrahedral meshes; however, there is some flexibility because of additional freedom in three dimensions. The application of this methodology to partitioning and load balancing holds a lot of promise and is currently being investigated. However, modifications to improve locality in the sense of a generalized frontal method requires a more detailed study of the relationships between more specialized constraints on walks and cache memory behavior.

## References

- [1] R. Biswas and R. C. Strawn, *Mesh quality control for multiply-refined tetrahedral grids*, Appl. Numer. Math., 20 (1996), pp. 337–348.
- [2] J. Gerlach and G. Heber, *Fundamentals of natural indexing for simplex finite elements in two and three dimensions*, RWCP Technical Report, TR 97-008, Tsukuba Research Center, Japan, 1997.
- [3] G. Heber, R. Biswas, and G. R. Gao, *Self-avoiding walks over two-dimensional adaptive unstructured grids*, NAS Technical Report, NAS-98-007, NASA Ames Research Center, Moffett Field, 1998.

